

Department of Defense
High Performance Computing Modernization Program



Supercomputing:

HPCMP, Performance Measures and Opportunities

Cray J. Henry
August 2004

Report Documentation Page			Form Approved OMB No. 0704-0188	
<p>Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p>				
1. REPORT DATE 01 FEB 2005	2. REPORT TYPE N/A	3. DATES COVERED -		
4. TITLE AND SUBTITLE Supercomputing:HPCMP, Performance Measures and Opportunities			5a. CONTRACT NUMBER	
			5b. GRANT NUMBER	
			5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)			5d. PROJECT NUMBER	
			5e. TASK NUMBER	
			5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Department of Defense			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSOR/MONITOR'S ACRONYM(S)	
			11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution unlimited				
13. SUPPLEMENTARY NOTES See also ADM00001742, HPEC-7 Volume 1, Proceedings of the Eighth Annual High Performance Embedded Computing (HPEC) Workshops, 28-30 September 2004 Volume 1., The original document contains color images.				
14. ABSTRACT				
15. SUBJECT TERMS				
16. SECURITY CLASSIFICATION OF: a. REPORT b. ABSTRACT c. THIS PAGE unclassified unclassified unclassified			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 52
19a. NAME OF RESPONSIBLE PERSON				



Presentation Outline

● What's New in the HPCMP

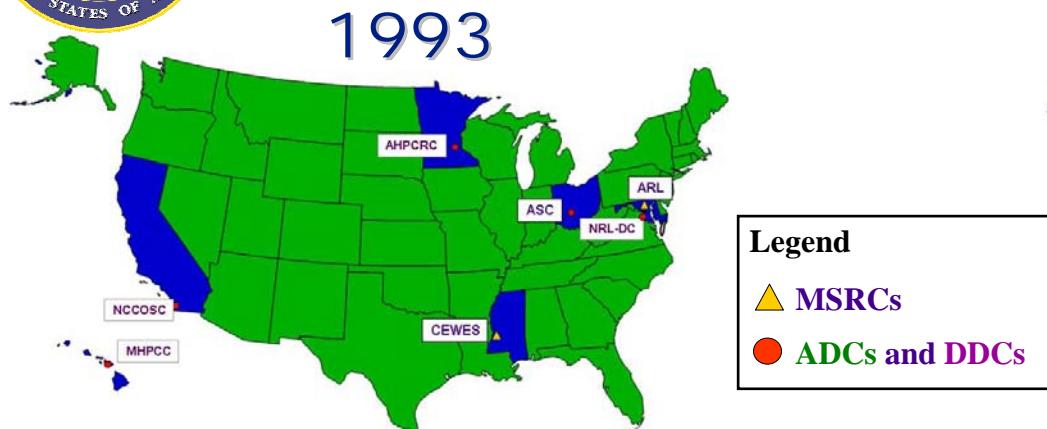
- New hardware
- HPC Software Application Institutes
- Capability Allocations
- Open Research Systems
- On-demand Computing

● Performance Measures - HPCMP

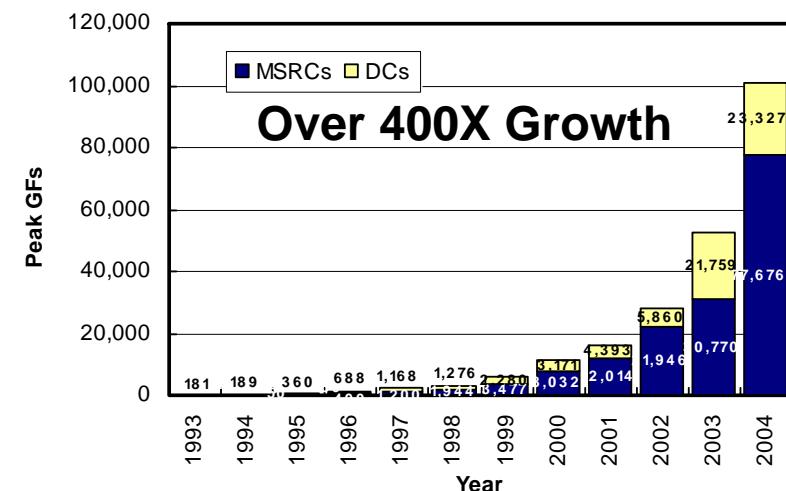
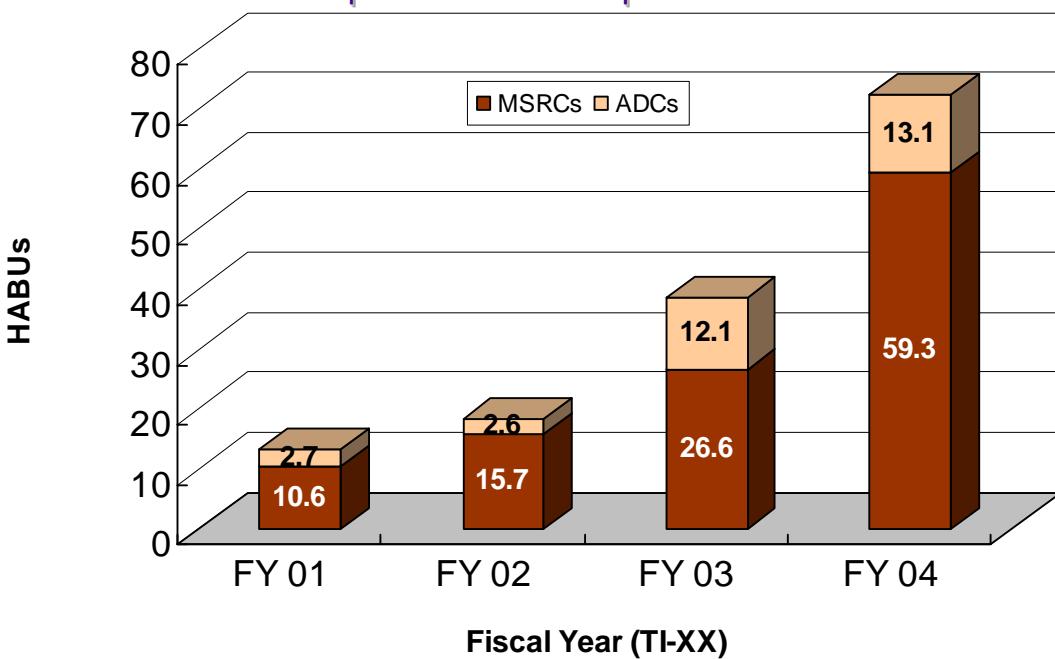
● Performance Measures – Challenges & Opportunities



HPCMP Centers



Total HPCMP End-of-Year Computational Capabilities





HPCMP Systems (MSRCs) 2004 HPEC Conference

Major Shared Resource Centers

FY 01 and earlier
FY 02
FY 03
FY 04

HPC Center	System	Processors
Army Research Laboratory (ARL)	IBM P3 SGI Origin 3800	1,280 PEs 256 PEs 512 PEs
	IBM P4	768 PEs 128 PEs
	Linux Networx Cluster LNX1 Xeon Cluster IBM Opteron Cluster SGI Altix Cluster	256 PEs 2,100 PEs 2,372 PEs 256 PEs
Aeronautical Systems Center (ASC)	Compaq SC-45 IBM P3 COMPAQ SC-40 SGI Origin 3900 SGI Origin 3900 IBM P4	836 PEs 528 PEs 64 PEs 2,048 PEs 128 PEs 32 PEs
Engineer Research and Development Center (ERDC)	Compaq SC-40 Compaq SC-45 SGI Origin 3800 Cray T3E SGI Origin 3900 Cray X1	512 PEs 512 PEs 512 PEs 1,888 PEs 1,024 PEs 64 PEs
Naval Oceanographic Office (NAVO)	IBM P4 SV1 IBM P4	1,408 PEs 64 PEs 3,456 PEs





HPCMP Systems (ADCs)



HPC Center	System	Processors	
Army High Performance Computing Center (AHCRC)	Cray T3E Cray X1, LC	1,088 PEs 128 PEs 64 PEs	FY 01 and earlier FY 02 FY 03
Arctic Region Supercomputing Center (ARSC)	Cray T3E Cray SV1 IBM P3 IBM Regatta P4 Cray X1	272 PEs 32 PEs 200 PEs 800 PEs 128 PEs	FY 04 upgrades
Maui High Performance Computing Center (MHPCC)	IBM P3 (2) IBM Netfinity Cluster IBM P4	736/320 PEs 512 PEs 320 PEs	
Space & Missile Defense Command (SMDC)	SGI Origins Cray SV-1 W.S. Cluster IBM e1300 Cluster Linux Cluster IBM Regatta P4	1,200 PEs 32 PEs 64 PEs 256 PEs 256 PEs 32 PEs	Why is the date important? Generally we see price-performance gains of ~ 1.68 (e.g., 2001 = 1 2002 = 1.68 x 2003 = 2.82 x 2004 = 4.74 x



HPCMP Dedicated Distributed Centers

Location	System	Description (Processors/Memory)
	Arnold Engineering Development Center (AEDC)	HP Superdome IBM Itanium Cluster IBM Regatta P4 Pentium Cluster
	Air Force Research Laboratory, Information Directorate (AFRL/IF)	Sky HPC-1
	Air Force Weather Agency (AFWA)	IBM Regatta P4 Heterogeneous HPC
	Aberdeen Test Center (ATC)	Powerwulf Powerwulf
	Fleet Numerical Meteorology and Oceanography Center (FNMOC)	SGI Origin3900 IBM Regatta P4
	Joint Forces Command (JFCOM)	Xeon Cluster

FY 04 new systems and/or upgrades



HPCMP Dedicated Distributed Centers

Location	System	Description (Processors/Memory)
 Naval Air Warfare Center, Aircraft Division (NAWCAD)	SGI Origin 2000 SGI Origin 3900	30 PEs 64 PEs
 Naval Research Laboratory-DC (NRL-DC)	SUN Sunfire 6800 Cray MTA SGI Altix SGI Origin 3000	32 PEs 40 PEs 128 PEs 128 PEs
 Redstone Technical Test Center (RTTC)	SGI Origin 3900	28 PEs
 Simulations & Analysis Facility (SIMAF)	SGI Origin 3900 Beowulf Cluster	24 PEs
 Space and Naval Warfare Systems Center-San Diego (SSCSD)	Linux Cluster IBM Regatta P4	128 PEs 128 PEs
 Whites Sands Missile Range (WSMR)	Linux Networx	64 PEs

FY 04 new systems and/or upgrades



Center POC's

Name	Org	Web URL	Contact Information
Brad Comes	HPCMO	http://www.hpcmo.hpc.mil	703-812-8205, bcomes@hpcmo.hpc.mil
Tom Kendall	ARL MSRC	http://www.arl.hpc.mil	410-278-9195 tkendall@arl.army.mil
Jeff Graham	ASC MSRC	http://www.asc.hpc.mil/	937-904-5135, Jeff.Graham@wpafb.af.mil
Chris Flynn	AFRL Rome DC	http://www.if.afrl.af.mil/tech/facilities/HPC/hpcf.html	315-330-3249, Christopher.Flynn@rl.af.mil
Dr. Lynn Parnell	SSCSD DC	http://www.spawar.navy.mil/sandiego/	619-553-1592, parnell@sscsd.hpc.mil
Maj Kevin Benedict	MHPCC DC	http://www.mhpcc.edu	808-874-1604, Kevin.Benedict@maui.afmc.af.mil



Disaster Recovery

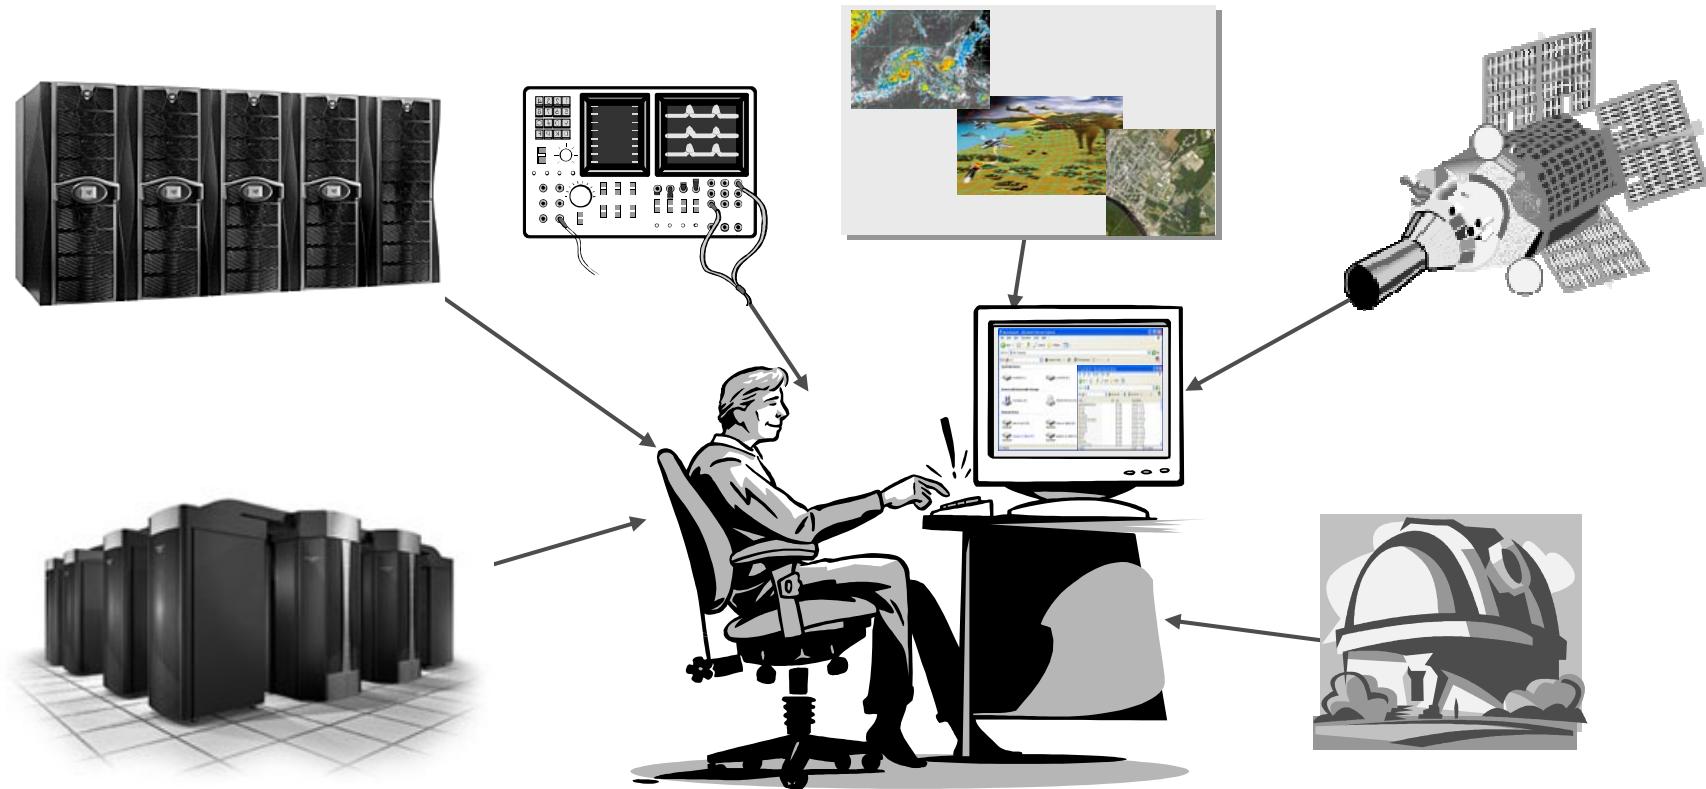
Retain third-copy of critical data at a hardened backup site so users can access their files from an alternate site in the event of disruption of their primary support site

- Status:
 - All MSRCs, MHPCC, and ARSC will have “off-site” third-copy backup storage for critical data
 - On-going initiative
- Working with centers to document the kinds of data that would need to be recovered
- Implementation to begin Q1 FY05



User Interface Toolkit

Provide an API-based toolkit to the user community and developers that facilitates the implementation of web-based interfaces to HPC

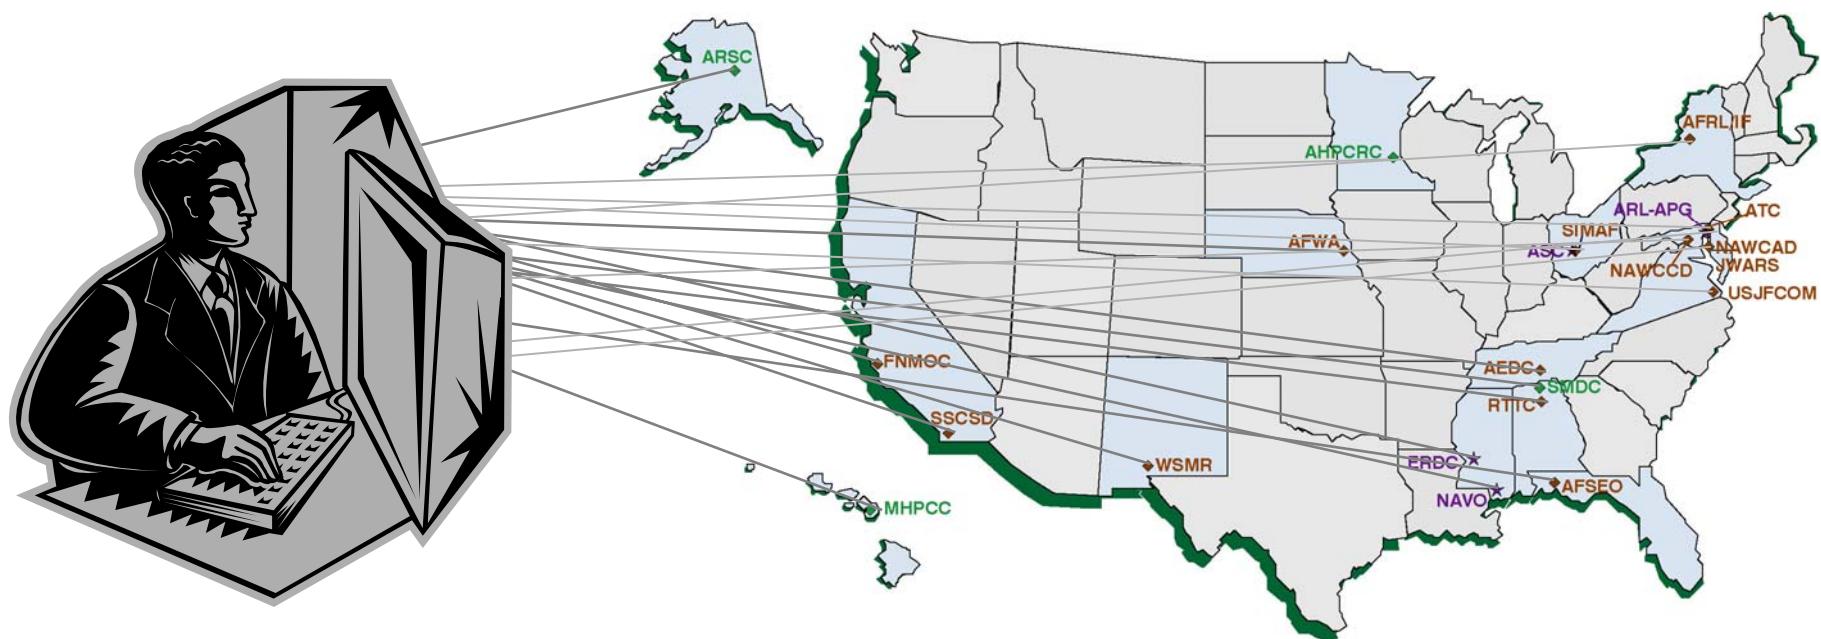


Facilitates Information Integration



Baseline Configuration

***Implement and Sustain a Common Set of Capabilities and Functions
Across the HPCMP Centers***



***Enables Users to Easily Move Between Centers Without the Requirement
to Learn and Adapt to Unique Configurations***



Software Applications Support

HPC Software Applications Institutes



- Lasting impact on services
- High value service programs

HPC Software Portfolios



- Tightly integrated software
- Address top DoD S&T and T&E problems

PET Partners



- Transfer of new technologies from universities
- On-site support
- Training

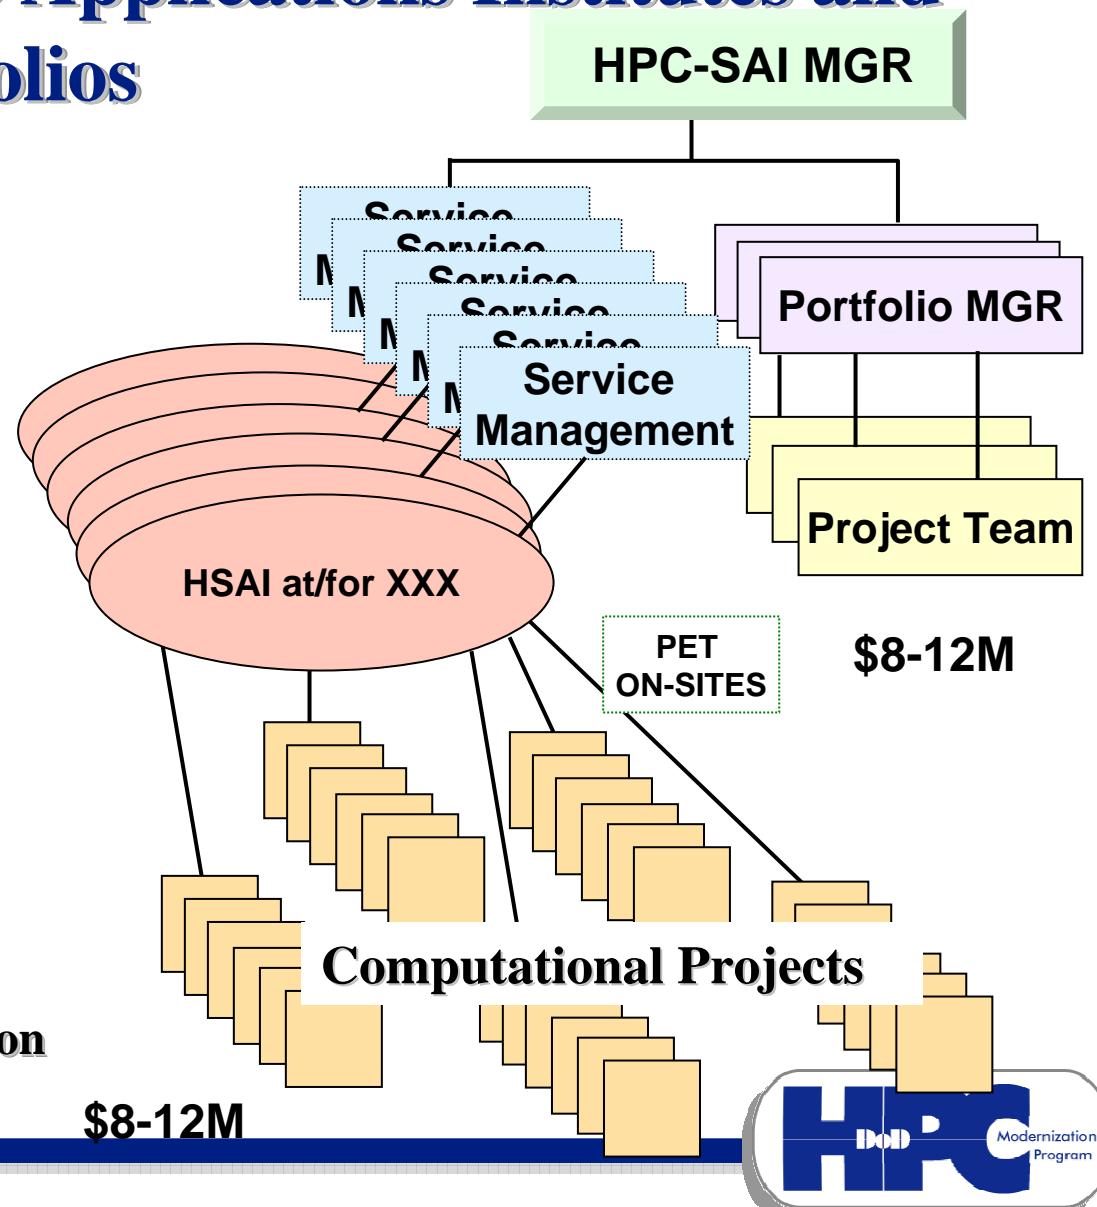


- Assure software intended use/user
- Protect software through source insertion



HPC Software Applications Institutes and Focused Portfolios

- 5–8 HPC Software (Applications) Institutes
 - HPCMP chartered
 - Service managed
 - 3–6 year duration
 - Ends with Transition to Local Support
 - \$0.5–3M annual funding for:
 - 3–12 computational and computer scientists
 - Support development of new and existing codes
 - Adjust local business practice to use science-based models & simulation
 - Integrated with PET

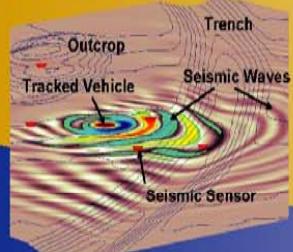




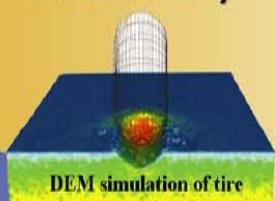
HPC Software Applications Institute (HSAI)

Institute for Maneuverability and Terrain Physics Simulation, ERDC

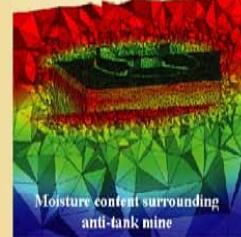
Acoustics/Seismics



Terrain Mobility



Countermine/IED/UXO

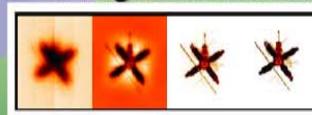


Institute for HPC Applications to Air Armament (IHAAA), AFSEO

Revolutionize Arming the Warfighter



Image Enhancement



Counterspace/Space-Based Surveillance



Biotechnology HSAI for Force Health Protection, MRMC

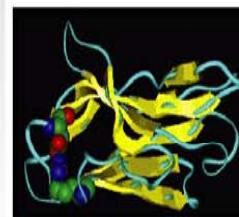
Force Health Protection

Experimentation

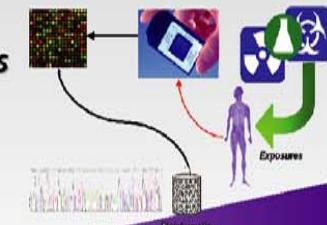


Simulation

Threat Detection/Diagnosis

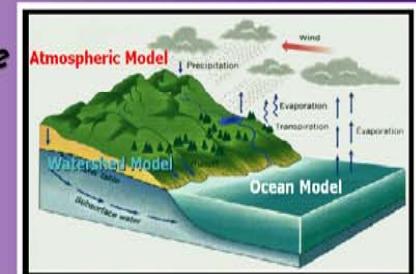


Computational Prediction of Protein Structure/Function

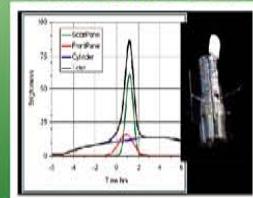


Battlespace Environments Institute, NRL

Earth and Space Modeling



Non-Imaging Space Object Identification and Data Fusion



Astrodynamics



HPC Software Applications Institute for Space Situation Awareness (SSA), AFRL



HPC Computational Fellowships

- Patterned after successful DOE fellowship program
- National Defense Science and Engineering Graduate Fellowship Program (NDSEG) chosen as vehicle for execution of fellowships
 - HPCMP added as fellowship sponsor along with Army, Navy, and Air Force
 - Computer **and computational** sciences added as possible discipline
- HPCMP **is sponsoring 11** fellows for 2004 and similar numbers each following year
- HPCMP fellows are strongly encouraged to develop close ties with DoD laboratories or test centers, including summer research projects
- User organizations have responded to DUSD (S&T) memo with fellowship POCs to select and interact with fellows



HPCMP Resource Allocation Policy

Capability Allocations

Goal: Support the top capability work

How:

- New TI-XX resources generally are implemented for a few months before the end of the current fiscal year without formal allocation
- Dedicate major fractions of large new systems to short-term, massive computations that generally cannot be addressed under normal shared resource operations for the first 2–3 months of life
- HPCMP issued call for short-term Capability Application Project (CAP) proposals
- Capability Application Projects will be implemented between October and December on large new systems each year
 - Proposals are required to show that the application efficiently used on the order of 1,000 processors or more and would solve a very difficult, important short-term computational problem



Status of Capability Applications Projects

- Call released to HPCMP community on 22 April 2004 with responses sent to HPCMPO by 1 June 2004
 - 21 proposals received across all large CTAs (CSM, CFD, CCM, CEA, and CWO)
- CAPs will be run on new 3,000 processor Power4+ at NAVO, 2,100 processor Xeon and 2,300 processor Opteron clusters at ARL
- CAPs will be run in two phases:
 - Exploratory phase designed to test scalability and efficiency of application codes to significant fractions of systems (5-15 projects on each system)
 - Production phase designed to accomplish significant capability work with efficient, scalable codes (1-3 projects on each system)
- Production phase of CAPs will be run after normal acceptance testing and pioneer work on these systems



“Open Research” Systems

- In response to customer demand: -- ~ 50% of Challenge Project leaders prefer to use an “open research” system
- “Open Research” systems concentrate on basic research allowing better separation of sensitive and non-sensitive information
 - minimal background check facilitating graduate student and foreign national access
- For FY05 the systems at ARSC will transition into an “open research” mode of operation
 - Eliminate the requirement for users of that system to have NACs
 - Customers would have to “certify” that their work is unclassified non-sensitive (e.g., open literature, basic research)
 - All other operational and security policies apply, such as all users of HPCMP resources must be valid DoD users assigned to a DoD computational project
 - Consistent with Uniform Use-Access Policy
- The account application process for “open research” centers or systems require certification by government program manager that computational work is cleared for open literature publication
 - Component of FY 2005 account request
- Operations on all other systems remain under current policies



On-demand (Interactive) Systems

- "Real-time" community has asked for "guaranteed" or on-demand service from shared resource centers
 - Request is aimed at ensuring quick response time from shared resource when system is being used interactively
 - Results needed now — can't wait
- Current policy requires that all Service/Agency work, be covered by an allocation
 - Note: "On-demand" system will have lower utilization but fast turn around
 - Service "valuation" of this service demonstrated by FY05 allocations — need sufficient allocation to dedicate a system to this mode of support
- Anticipating the Services/Agencies will allocate sufficient time to dedicate one 256 processor cluster at ARL



On-Demand Application --Distributed Interactive HPC Testbed

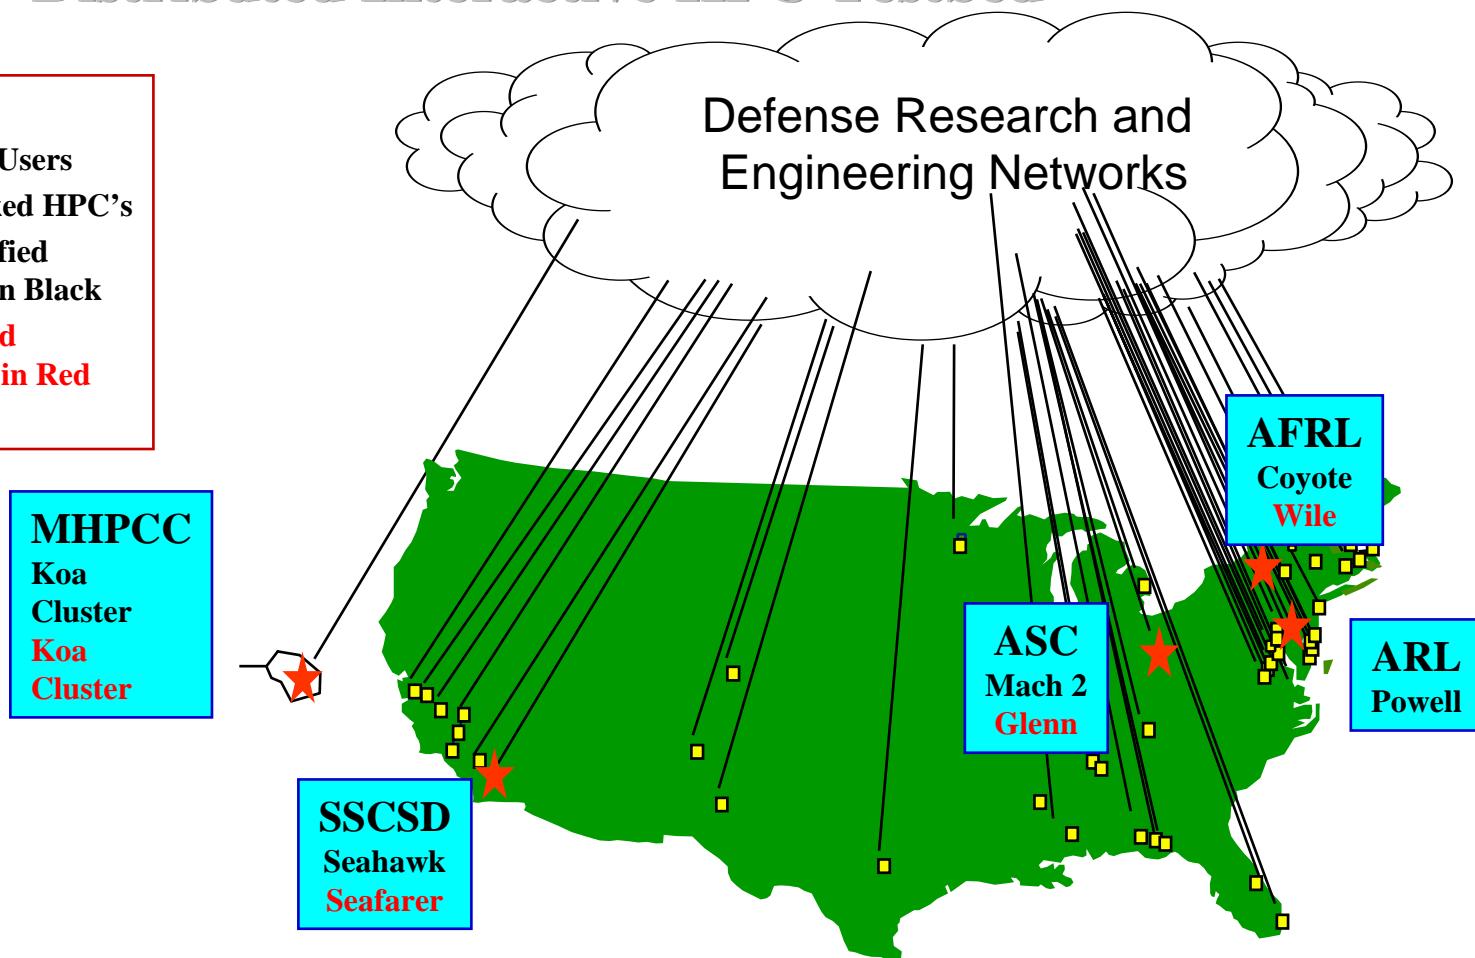
- **Goal:** Assess the potential value and cost of providing greater interactive access to HPC resources to the DoD RDT&E community and its contractors.
- **Means:** Provide both unclassified and classified distributed HPC resources to the DoD HPC community in FY05 for interactive experimentation exploring new applications and system configurations



Distributed Interactive HPC Testbed

Legend

- [Yellow Square] Remote Users
- [Red Star] Networked HPC's
- Unclassified System in Black
- Classified Systems in Red



- ✓ Distributed HPC's
- ✓ Accessed by authorized users anywhere on the DREN and Internet
- ✓ Interactive and time critical problems



-- Technical Challenges

- Low latency support for interactive and real-time applications—proper HPC configuration?
- Cohabitation of interactive and batch jobs?
- Web-based access to network of HPC's with enhanced usability
- Consistency with HPCMP approved secure environment using DREN and SDREN
- Information management system supporting distributed HPC applications
- Demonstrating new C4ISR applications of HPC
- Expanding FMS use beyond Joint experimentation to include training and mission rehearsal



--Interactive Parallel MATLAB at

- **Objectives:** to provide SIP users with a High Productivity Interactive Parallel MATLAB environment (it will provide the user-friendly MATLAB high-level language syntax plus the computational power of the interactive HPCs)
- **To allow interactive experiments** for demanding SIP problems: problems that take too long to finish on a single Workstation, or that require more memory than what is available on a single computer, or systems with both constraints in which users' research may benefit by an interactive modus-operandi.
- **Approach:** to use MatlabMPI or other Parallel MATLAB viable approaches to deliver parallel execution but keeping the familiar MATLAB interactive environment
- It may serve as a vehicle to collect experimental data about **productivity issues:** are SIP users really more productive on such an *Interactive HPC MATLAB platform?* (versus the traditional batch oriented HPCs)



DIHT High Performance Computers

2004 HPEC Conference

Site	Computer	Memory and I/O	Online
ARL MSRC Aberdeen, MD	Unclass- Powell: 128 node Dual 3.06MHz Xeon Cluster	2 GB DRAM and 64 GB disk/node, Myrinet & GigEnet/ 100MB Backplane	Est. 10/04 w/batch; 4/05 share with batch,
ASC MSRC Dayton, OH	Unclass- Mach2: 24 node Dual 2.66 GHz Xeon, Linux Class -Glenn: 128 node dual Xeon, Linux	4 GB DRAM and 80 GB disk/node , dual GigEnet 4 GB DRAM and local disks	Est. 10/04 Est. Spring/05
AFRL Rome, NY	Unclass- Coyote: 26 node Dual 3.06GHz Xeon, Linux Class - Wile:14 node Dual 2.66/3.06 GHz Xeon, Linux	6 GB DRAM and 400 GB disk/node, dual GigEnet 6 GB DRAM and 200 GB disk/node, dual GigEnet	Yes Est. 12/04
SSCSD San Diego, CA	Unclass- Seahawk: 16 node 1.3GHz Itanium2, Linux Class - Seafarer: 24 node Dual 3.06 GHz	2 GB DRAM and 36 GB disk/node, dual GigEnet 4 GB DRAM and 80 GB disk/node, dual GigEnet	Est. 12/04 Yes (U) til 3/05
MHPCC Maui, HI	Unclass/ Class - Koa: 128 node dual Xeon, Linux (system moves between environments)	4 GB DRAM and 80 GB disk/node, shared file system, dual GigEnet	Yes





Key Technical Users

2004 HPEC Conference

Name	Program	Contact Information
Dr. Richard Linderman	HPC for Information Management	315-330-2208, Richard.Linderman@rl.af.mil
Dr. Bob Lucas	USJFCOM J9	310-448-9449, rflucas@isi.edu
Dr. Stan Ahalt	PET- SIP CTP	614-292-9524, ahalt@osc.edu
Dr. Juan Carlos Chaves	Interactive Parallel MATLAB	410-278-7519, jchaves@arl.army.mil
Dr. Dave Pratt	SBA Force transformations	407-243-3308, David.R.Pratt@saic.com
Rob Ehret Bill McQuay	Grid-based Collaboration	937-904-9017, Robert.Ehret@sensors.wpafb.af.mil 937-904-9214, William.Quay@sensors.wpafb.af.mil
Dr. John Nehrbass	Web enabled HPC	937-904-5139, John.Nehrbass@wpafb.af.mil
Dr. Keith Bromley	Signal Image Processing	619-553-2535, bromley@spawar.navy.mil
Dr. George Ramseyer	Hyperspectral Image Exploitation	315-330-3492, George.Ramseyer@rl.af.mil
Richard Pei	Interactive Electromagnetics Sim	732-532-0365, Richard.Pei@us.army.mil
Dr. Ed Zelnio	3-D SAR Radar Imagery	937-255-4949 ext.4214, Ed_Zelnio@mbvlab.wpafb.af.mil
John Rooks	Swathbuckler SAR Radar Imagery	315-330-2618, John.Rooks@rl.af.mil

Department of Defense

High Performance Computing Modernization Program

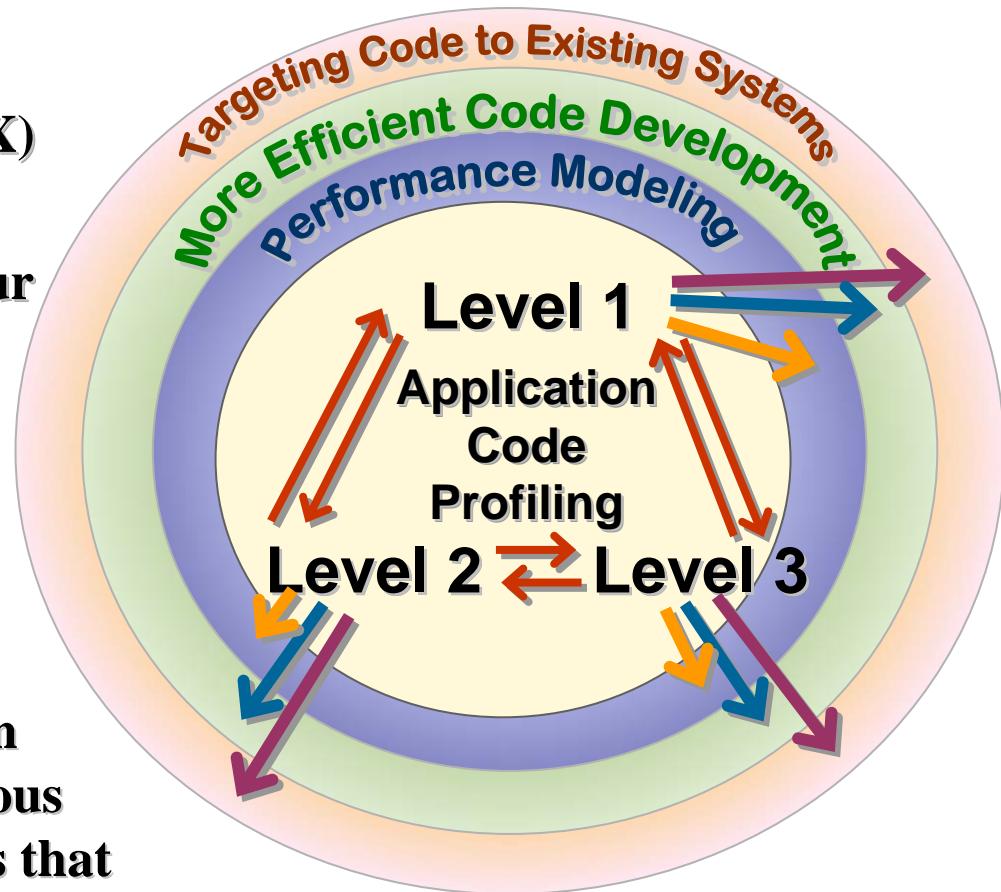
HPCMP Benchmarking and Performance Modeling Activities

<http://www.hpcmo.hpc.mil>



Performance Measurement Goals

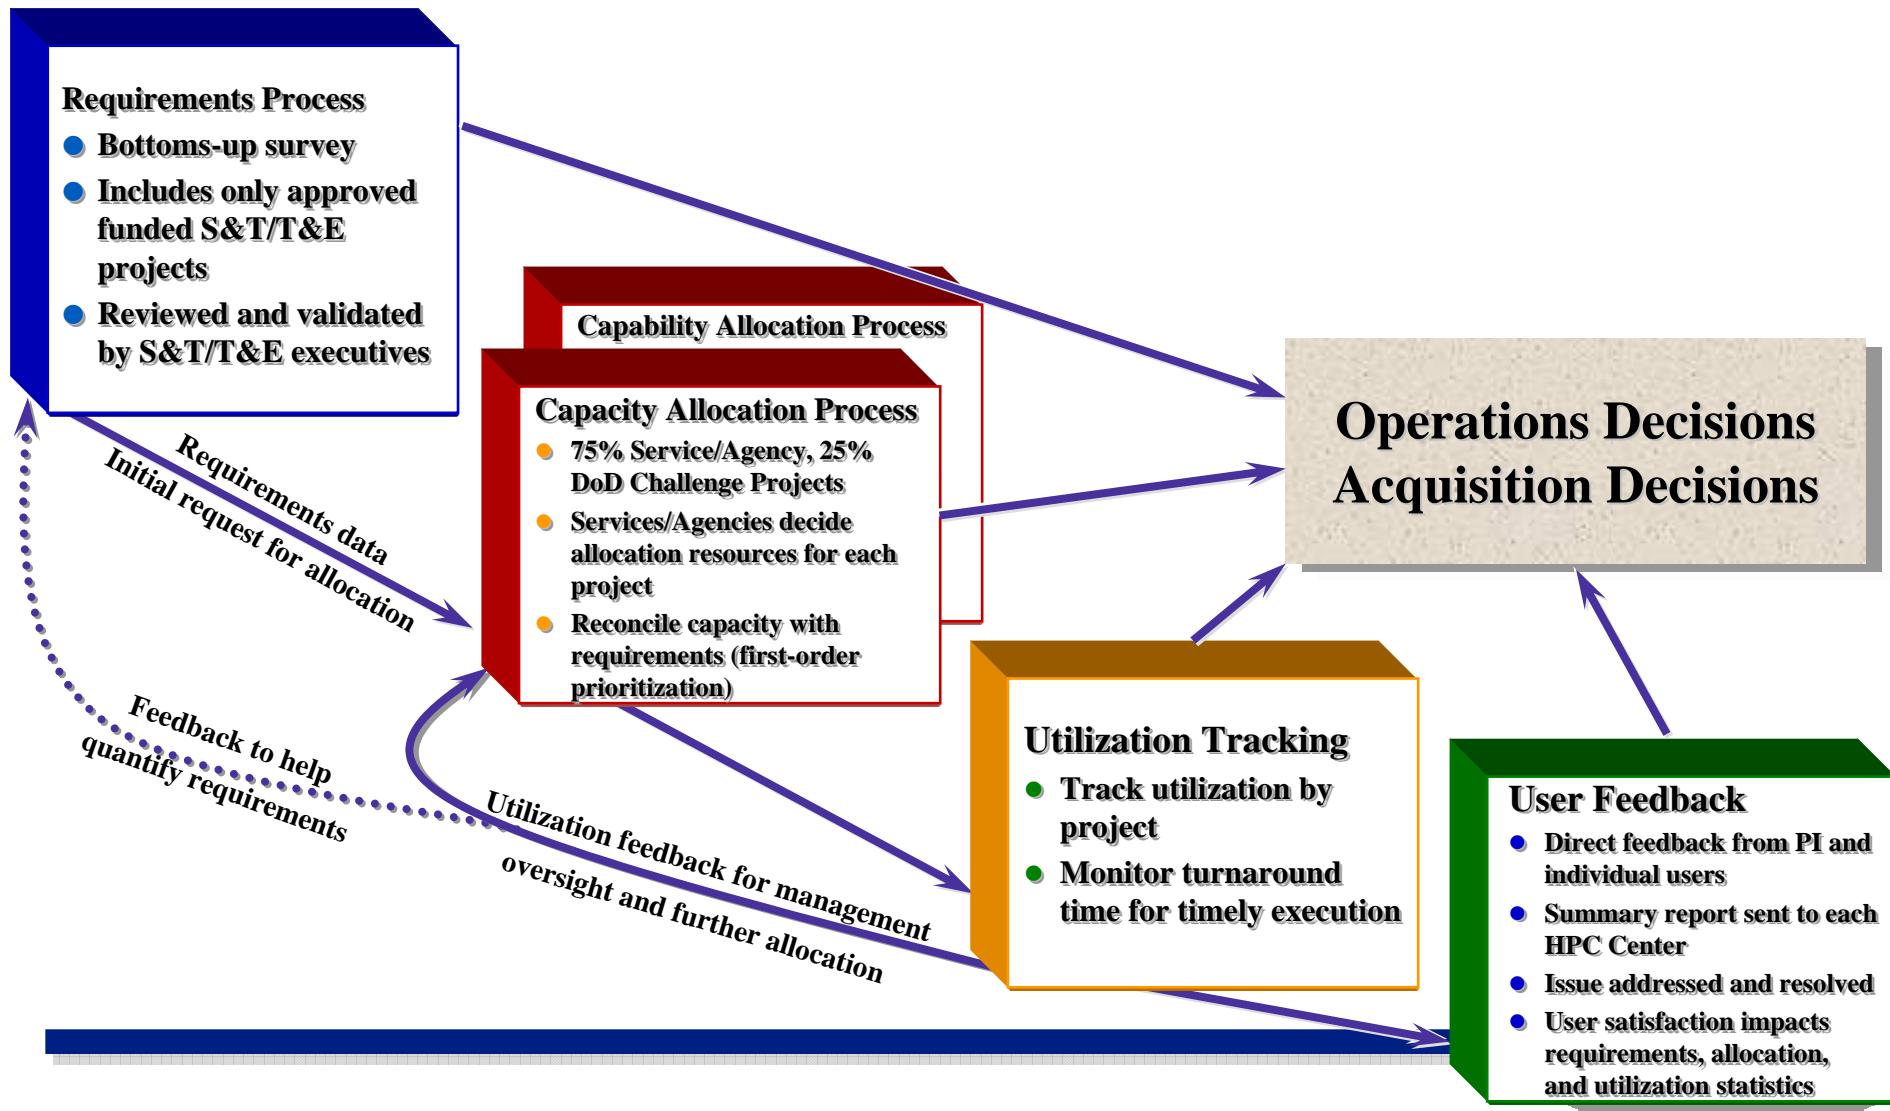
- Provide Quantitative measures to support selection of computers in annual procurement process (TI-XX)
 - Develop an understanding of our key application codes for the purpose of guiding code developers and users toward more efficient applications and machine assignments
 - Replace the current application benchmark suite with a judicious choice of synthetic benchmarks that could be used to predict performance of any HPC architecture on the program's key applications





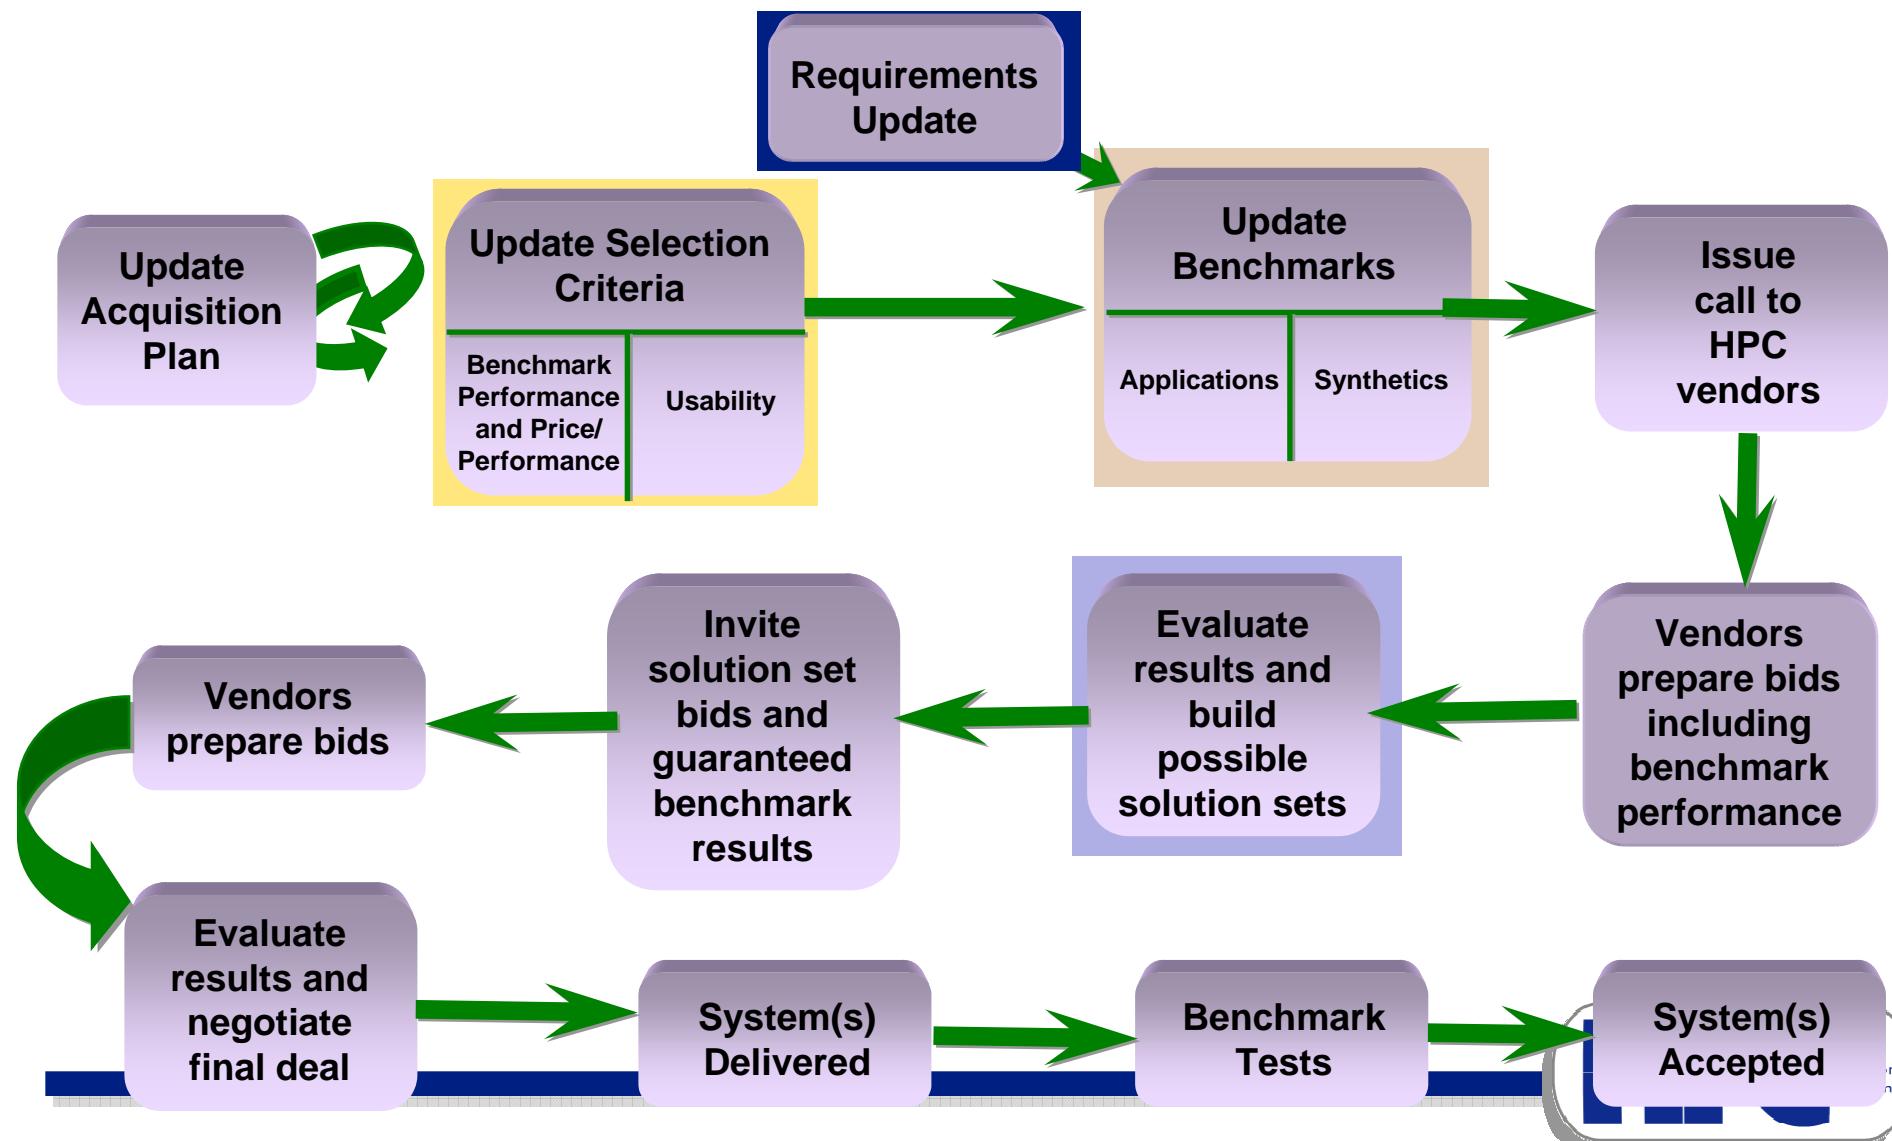
Resource Management

— Integrated Requirements/Allocation/Utilization Process





Technology Insertion (TI) Flow Chart





Types of Benchmark Codes

● Synthetic codes

- Basic hardware and system performance tests
- Meant to determine expected future performance
- Scalable, quantitative synthetic tests will be used for scoring and others will be used as system performance checks by Usability Team

● Application codes

- Actual application codes as determined by requirements and usage
- Meant to indicate current performance



Percentage of Unclassified Non-Real-Time Requirements, Usage, and Allocations

CTA	Requirements Percentage <i>FY [2002] (2003) {2004}</i>	Usage Percentage <i>FY 2002 {2003}</i>	Allocation Percentage <i>FY 2003 {2004}</i>	Average (25% FY 2004 Req, 25% FY 2003 Usage, 50% FY 2004 Alloc)
				<i>FY [2002] (2003) {2004}</i>
CFD	[35.5%] (36.9%) {38.6%}	48.3% {37.2%}	40.7% {44.4%}	[43.3%] (41.6%) {41.2%}
CCM	[15.5%] (18.6%) {16.2%}	16.4% {21.2%}	14.2% {12.6%}	[14.2%] (15.9%) {15.7%}
CWO	[21.9%] (19.2%) {20.8%}	21.3% {23.1%}	21.9% {17.6%}	[23.3%] (21.1%) {19.8%}
CEA	[4.1%] (4.0%) {4.8%}	5.1% {4.8%}	8.2% {6.6%}	[4.9%] (6.4%) {5.7%}
CSM	[11.4%] (11.8%) {11.7%}	3.5% {7.5%}	9.6% {11.0%}	[8.3%] (8.6%) {10.3%}
EQM	[3.0%] (3.2%) {2.1%}	0.6% {1.6%}	4.0% {3.1%}	[2.3%] (3.0%) {2.4%}
SIP	[1.0%] (1.4%) {1.4%}	1.2% {1.1%}	0.2% {0.4%}	[0.4%] (0.7%) {0.8%}
CEN	[0.5%] (0.4%) {0.6%}	1.3% {1.2%}	0.1% {1.2%}	[1.4%] (0.5%) {1.1%}
IMT	[2.9%] (0.8%) {0.8%}	2.1% {0.7%}	0.7% {1.9%}	[0.9%] (1.1%) {1.3%}
Other	[1.3%] (1.2%) {0.2%}	0.1% {0.8%}	0.2% {0.7%}	[0.4%] (0.4%) {0.6%}
FMS	[2.9%] (2.6%) {2.9%}	0.2% {0.8%}	0.2% {0.4%}	[0.7%] (0.8%) {1.1%}



TI-05 Application Benchmark Codes

- Aero – Aeroelasticity CFD code (single test case)
(Fortran, serial vector, 15,000 lines of code)
- AVUS (Cobalt-60) – Turbulent flow CFD code
(Fortran, MPI, 19,000 lines of code)
- GAMESS – Quantum chemistry code
(Fortran, MPI, 330,000 lines of code)
- HYCOM – Ocean circulation modeling code
(Fortran, MPI, 31,000 lines of code)
- OOCore – Out-of-core solver
(Fortran, MPI, 39,000 lines of code)
- RFCTH2 – Shock physics code
(~43% Fortran/~57% C, MPI, 436,000 lines of code)
- WRF – Multi-Agency mesoscale atmospheric modeling code (single test case)
(Fortran and C, MPI, 100,000 lines of code)
- Overflow-2 – CFD code originally developed by NASA
(Fortran 90, MPI, 83,000 lines of code)



TI-04 Benchmark Weights

CTA	Benchmark	Size	Unclassified %	Classified %
CSM	RF-CTH	Standard	a%	A%
CSM+CFD	RF-CTH	Large	b%	B%
CFD	Cobalt60	Standard	c%	C%
CFD	Cobalt60	Large	d%	D%
CFD	Aero	Standard	e%	E%
CEA+SIP	OOCore	Standard	f%	F%
CEA+SIP	OOCore	Large	g%	G%
CCM+CEN	GAMESS	Standard	h%	H%
CCM+CEN	GAMESS	Large	i%	I%
CCM	NAMD	Standard	j%	J%
CCM	NAMD	Large	k%	K%
CWO	HYCOM	Standard	l%	L%
CWO	HYCOM	Large	m%	M%
Total			100.00%	100.00%



Emphasis on Performance

- Establish a DoD standard benchmark time for each application benchmark case
 - NAVO IBM Regatta P4 (**Marcellus**) chosen as standard DoD system for TI-04 (**Initially IBM SP3 – HABU**)
- Benchmark timings (at least three on each test case) are requested for systems that meet or beat the DoD standard benchmark times by at least a factor of two (preferably up to four)
- Benchmark timings may be extrapolated provided they are guaranteed, but at least one actual timing on the offered or closely related system must be provided



CTH Standard

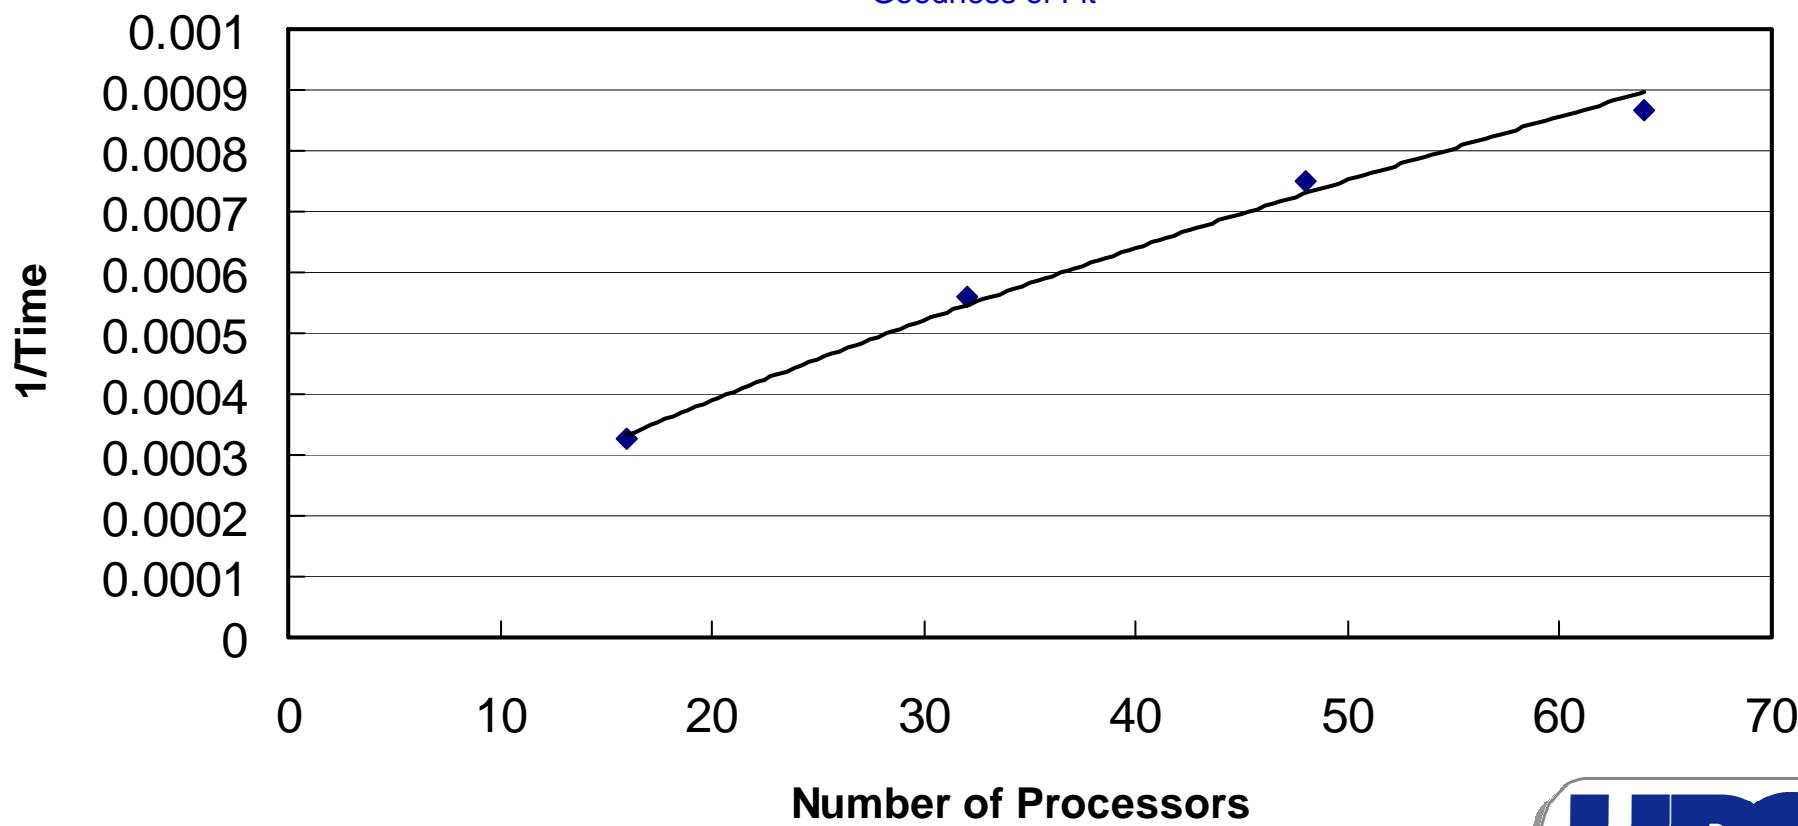
NAVO IBM SP P3 — 1288 Processors

“Slope”

$$y = 4.57590E-05x^{7.15387E-01}$$

$$R^2 = 9.94381E-01$$

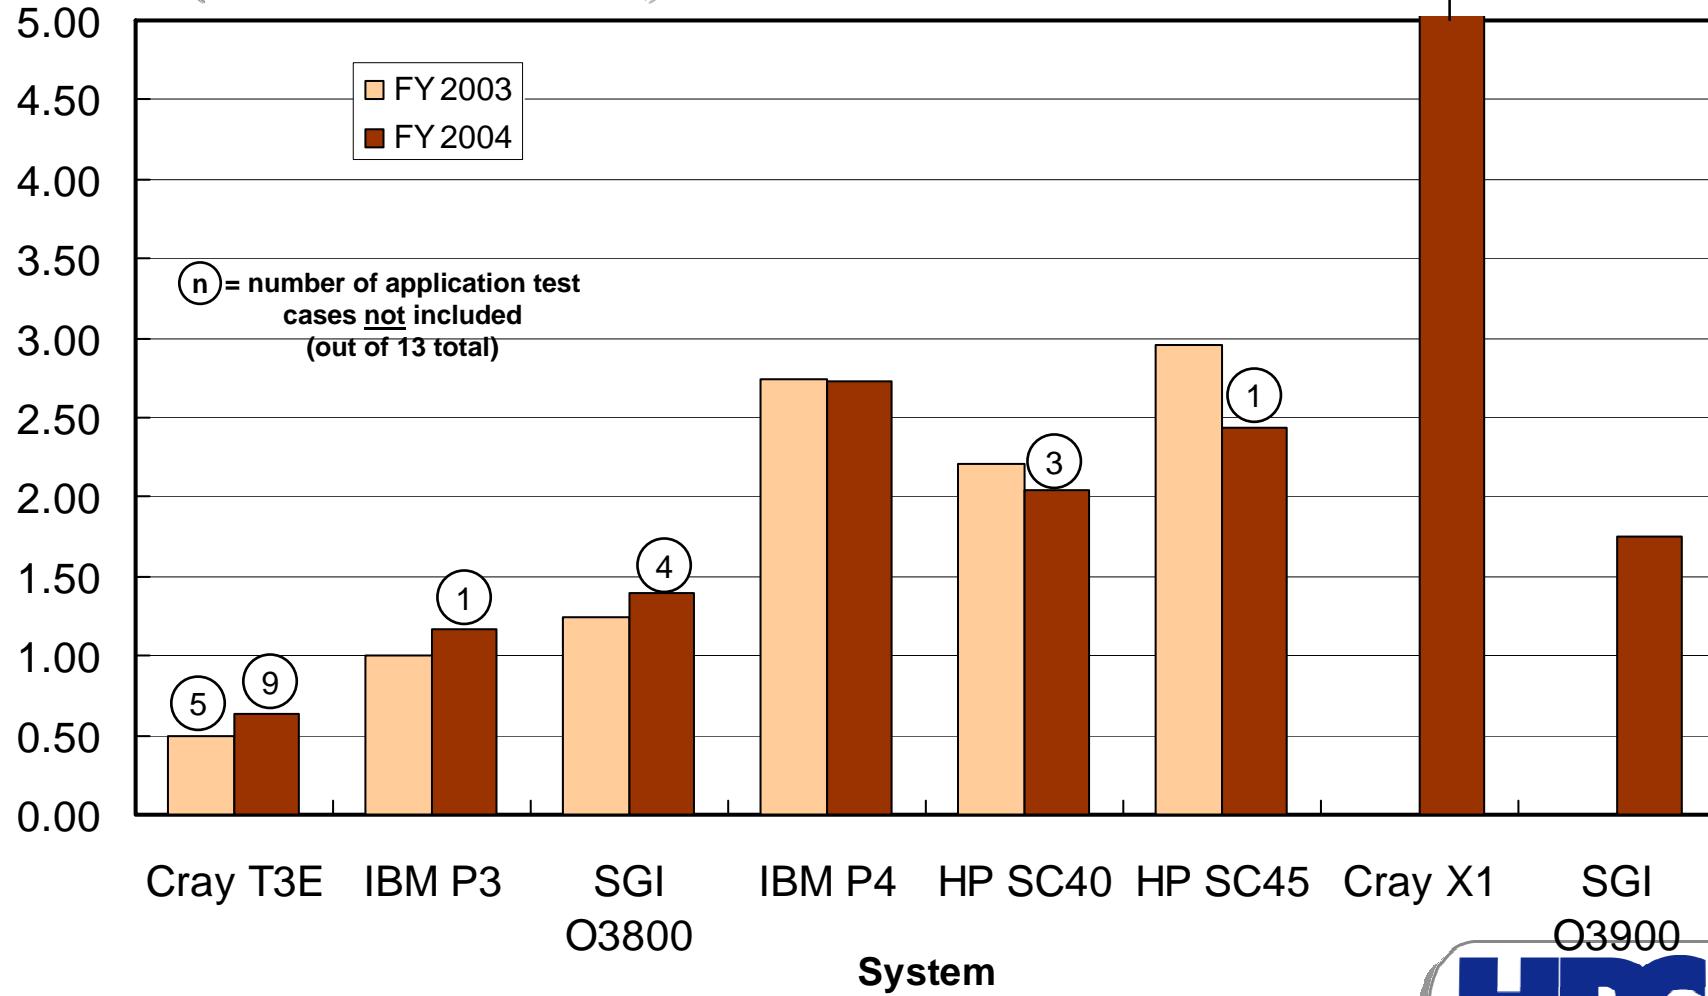
x = Number of Processors
 y = 1/Time





HPCMP System Performance (Unclassified)

Normalized Habu Equivalents





How the Optimizer Works: 2004 HPEC Conference

Problem Description

KNOWN

Prices

Application Score
Matrix

\$	
\$	
\$	
\$	
\$	
\$	
\$	
\$	

Machines

Application Test Case Codes								
S	S	S	S	S	S	S	S	S
S	S	S	S	S	S	S	S	S
S	S	S	S	S	S	S	S	S
S	S	S	S	S	S	S	S	S
S	S	S	S	S	S	S	S	S
S	S	S	S	S	S	S	S	S
S	S	S	S	S	S	S	S	S
S	S	S	S	S	S	S	S	S

Budget
Limits

Overall Desired

\$
\$

Workload Distribution

%	%	%	%	%	%	%	%
%	%	%	%	%	%	%	%

Application Test Case Codes

UNKNOWN

Optimal
Quantity
Set

#
#
#
#
#
#
#
#

Application Test Case Codes								
%	%	%	%	%	%	%	%	%
%	%	%	%	%	%	%	%	%
%	%	%	%	%	%	%	%	%
%	%	%	%	%	%	%	%	%
%	%	%	%	%	%	%	%	%
%	%	%	%	%	%	%	%	%
%	%	%	%	%	%	%	%	%
%	%	%	%	%	%	%	%	%

Optimize Total Price/Performance



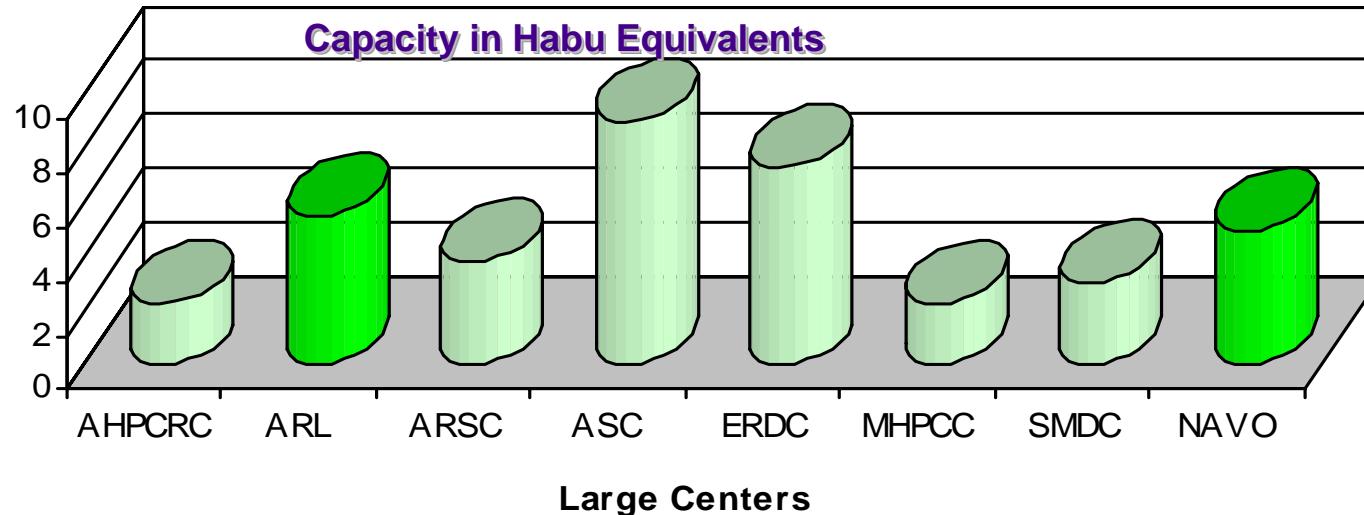
Price Performance Based Solutions

System	Total # Proc	Opt # 1	Opt # 2	Opt # 3	Opt # 4
A	64	1	1	0	0
B	188	0	2	3	0
C	128	0	0	0	4
C	256	0	2	4	0
D	256	15	0	0	12
D	512	0	4	1	1
E	256	1	1	3	0
Performance / Life Cycle		3.03	3.02	2.97	2.95

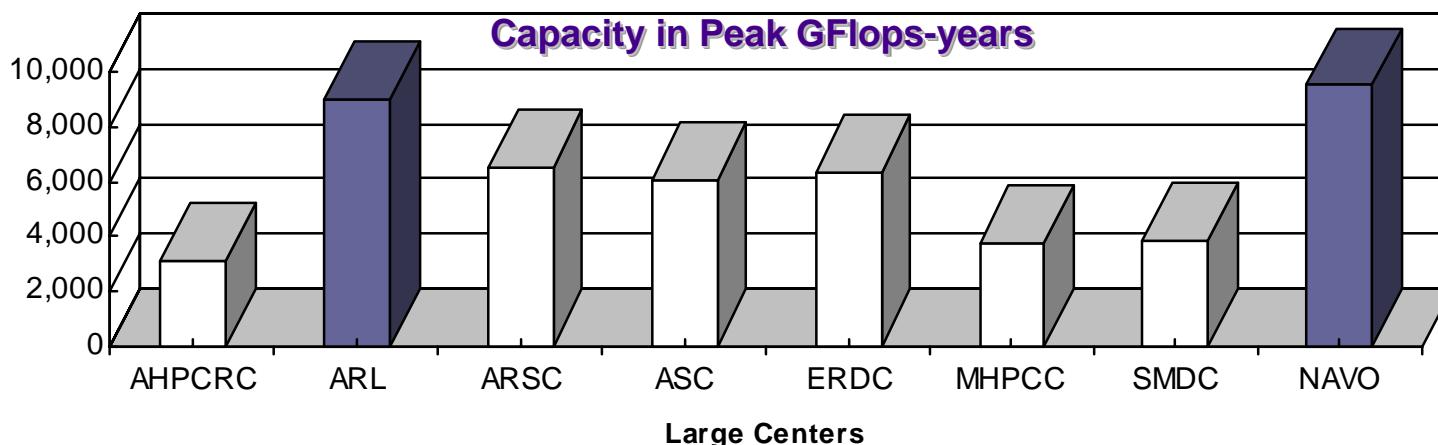
The optimizer produces a list of system solutions in rank order based upon Performance / Life Cycle Cost



Capturing True Performance Benchmarks



Top 500 or
Peak GFlops
is not a
Measure
of Real
Performance

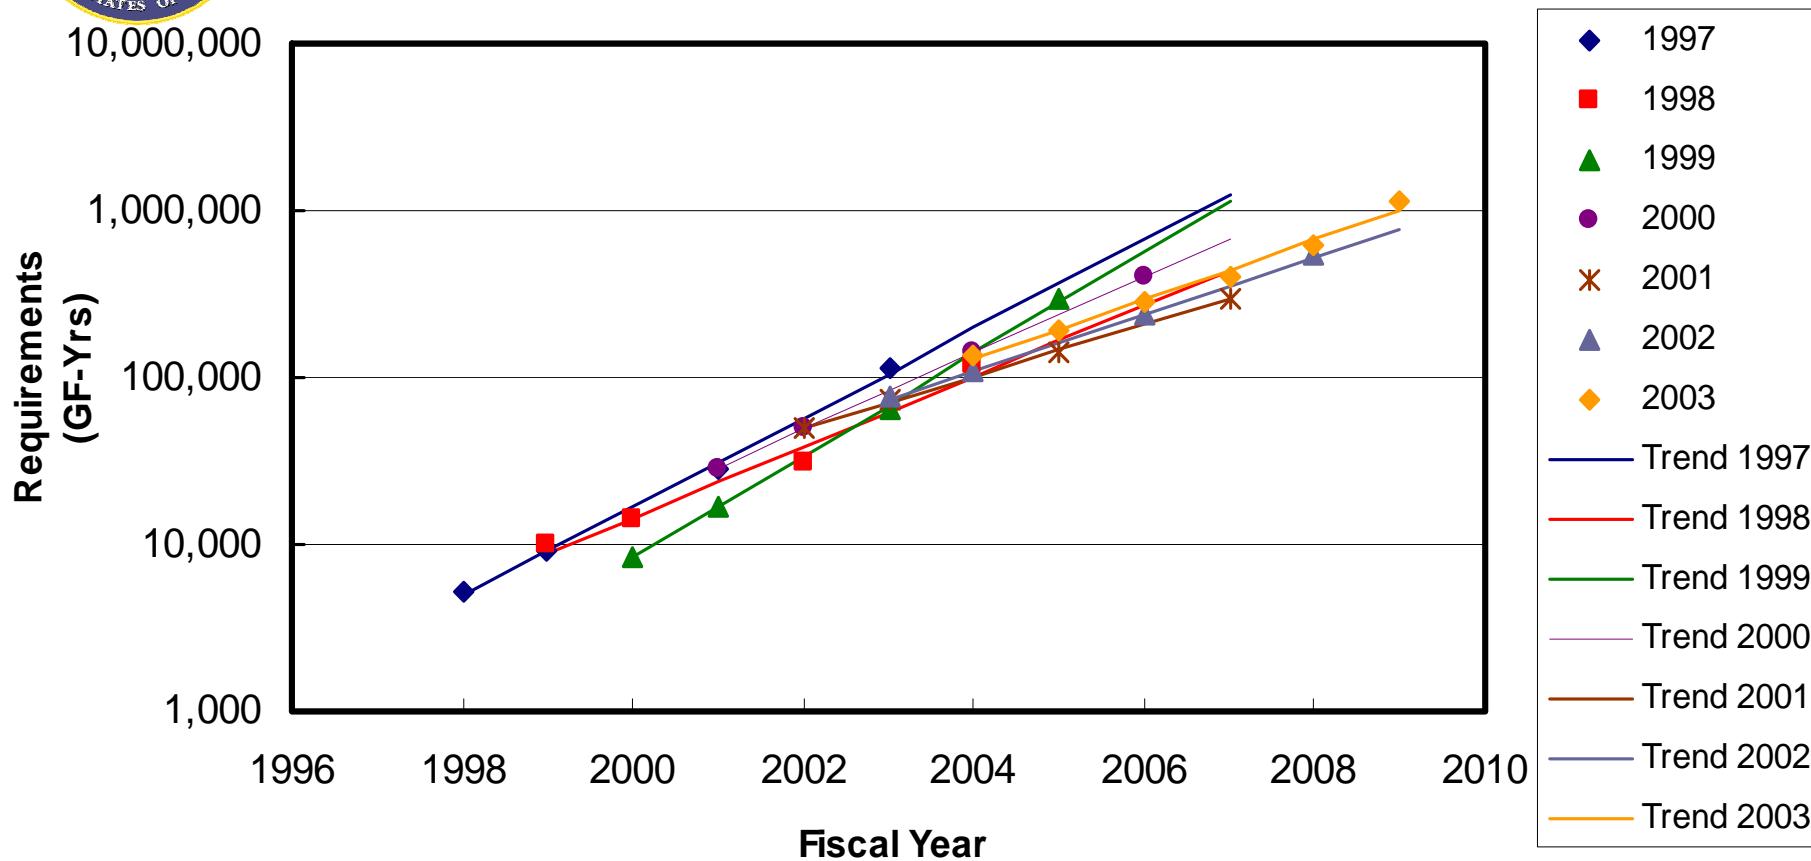


At the end
of TI-03



Requirement Trends

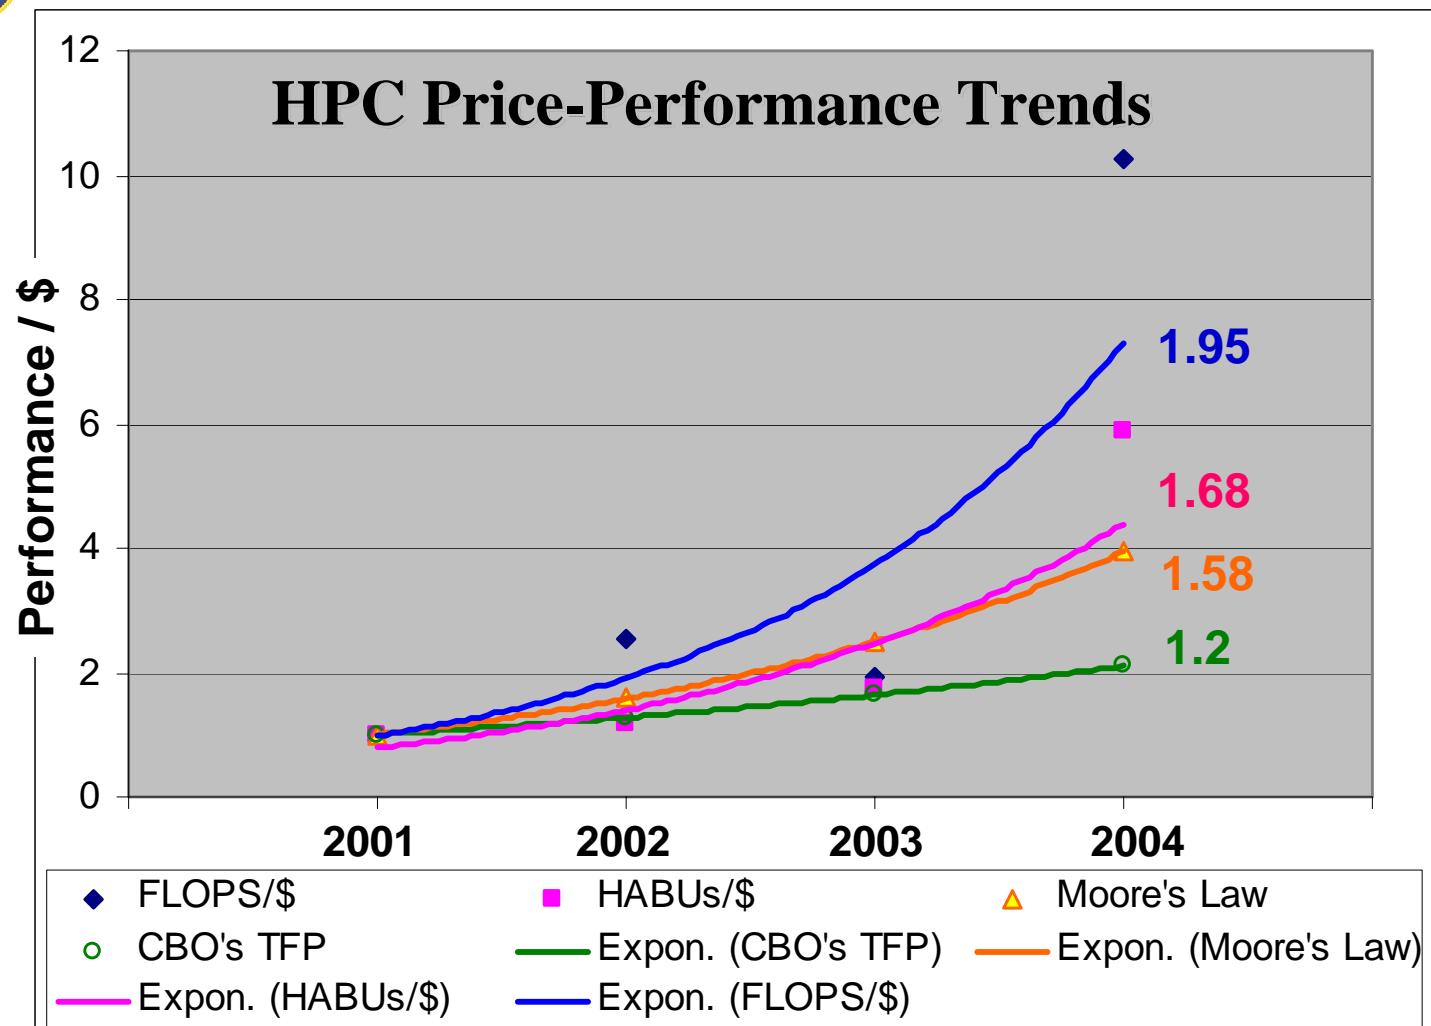
2004 HPEC Conference



The slope of this semi-log plot for the entire set of data equates to a constant factor of (1.76 ± 0.26) , although the slopes for the last two years have been 1.42 and 1.48, respectively.



Supercomputer Price-Performance Trends



Department of Defense

High Performance Computing Modernization Program

HPCMP Benchmarking and Performance Modeling

Challenges & Opportunities



Benchmarks

Today

Dedicated Applications

- 80% weight
- Real codes
- Representative data sets

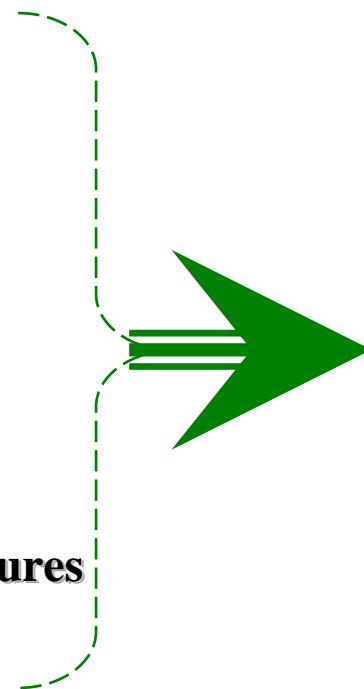
Synthetic Benchmarks

- 20% weight
- Future look
- Focus on key machine features

Tomorrow

Synthetic Benchmarks

- 100% weight
- Coordinated to application “signature”
- Performance on real codes accurately predicted from synthetic benchmark results
- Supported by genuine “signature” databases



Next 1–2 years key — must prove that synthetics benchmarks and application “signatures” can be coordinated



How -- Application Code Profiling Plan

- Began at behest of HPC User Forum in partnership with NSA
- Has evolved to multi-year plan -- how key application codes perform on HPC systems
 - Maximizing use of current HPC resources
 - Predicting performance of future HPC resources
- Performers include
 - Programming Environment and Training (PET) partners
 - Performance Modeling and Characterization Laboratory (PMaC) at SDSC
 - Computational Science and Engineering Group at ERDC
 - Instrumental, Inc.
- Research and production activities include
 - Profiling key DoD application codes at several different levels
 - Characterizing HPC systems with a set of system probes (synthetic benchmarks)
 - Predicting HPC system performance based on application profiles
 - Determining a minimal set of HPC system attributes necessary to model performance
 - Constructing the appropriate set of synthetic benchmarks to accurately model the HPCMP computational workload to use in system acquisitions



Support for TI-05 (Scope and Schedule)

- **Level 3 application code profiling**
 - **Eight application codes – 14 unique test cases**
 - **Each test case to be run at 3 different processor counts**
- **Predictions for existing systems**
 - **21 systems at 7 centers (some overlap possible in predictions)**
 - **Benchmarking POCs identified for each center**
 - **Goal: benchmarking results and predictions complete by Dec 2004**
- **Predictions for offered systems**
 - **Goal: benchmarking results finalized by 19 November 2004; all predictions completed by 31 December 2004**
- **Sensitivity Analysis**
 - **Goal: Determine how accurate a prediction do we need.**



Should We Do Uncertainty Analysis?





Performance Prediction Uncertainty Analysis

- Overall goal: Understand and accurately estimate uncertainties in performance predictions
- Determine functional form of performance prediction equations and develop uncertainty equation
- Determine uncertainties in underlying measured values from system probes and application profiling and use uncertainty equation to estimate uncertainties
- Compare results of performance prediction to measured timings and uncertainties of these results to predicted uncertainties
- Assess uncertainties in measured timings and determine whether acceptable agreement is obtained
- Eventual goal: propagate uncertainties in performance prediction to determine uncertainties in acquisition scoring



Performance Modeling Uncertainty Analysis

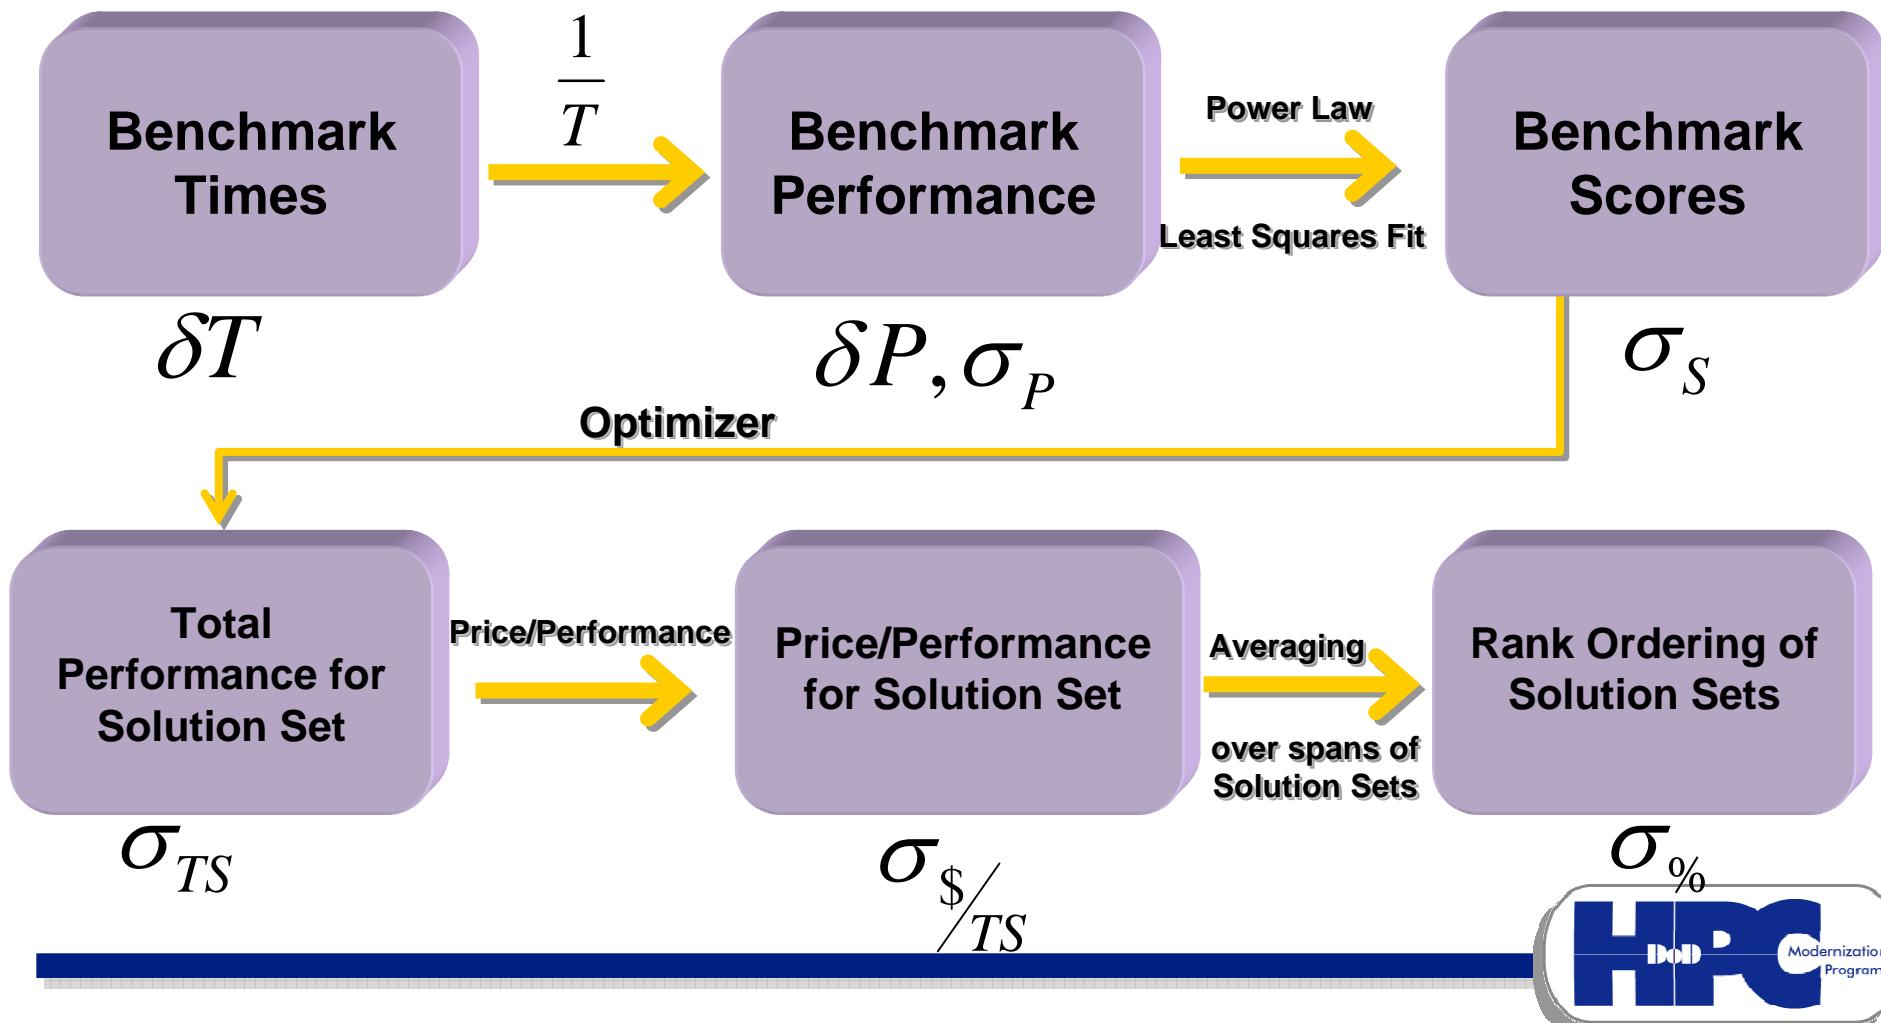
- Assumption: Uncertainties in measured performance values can be treated as uncertainties in measurements of physical quantities
- For small, random uncertainties in measured values x, y, z, \dots , the uncertainty in a calculated function $q(x, y, z \dots)$ can be expressed as:

$$\delta q = \sqrt{\left(\frac{\partial q}{\partial x} \delta x \right)^2 + \dots + \left(\frac{\partial q}{\partial z} \delta z \right)^2}$$

- Systematic errors need careful consideration since they cannot be calculated analytically

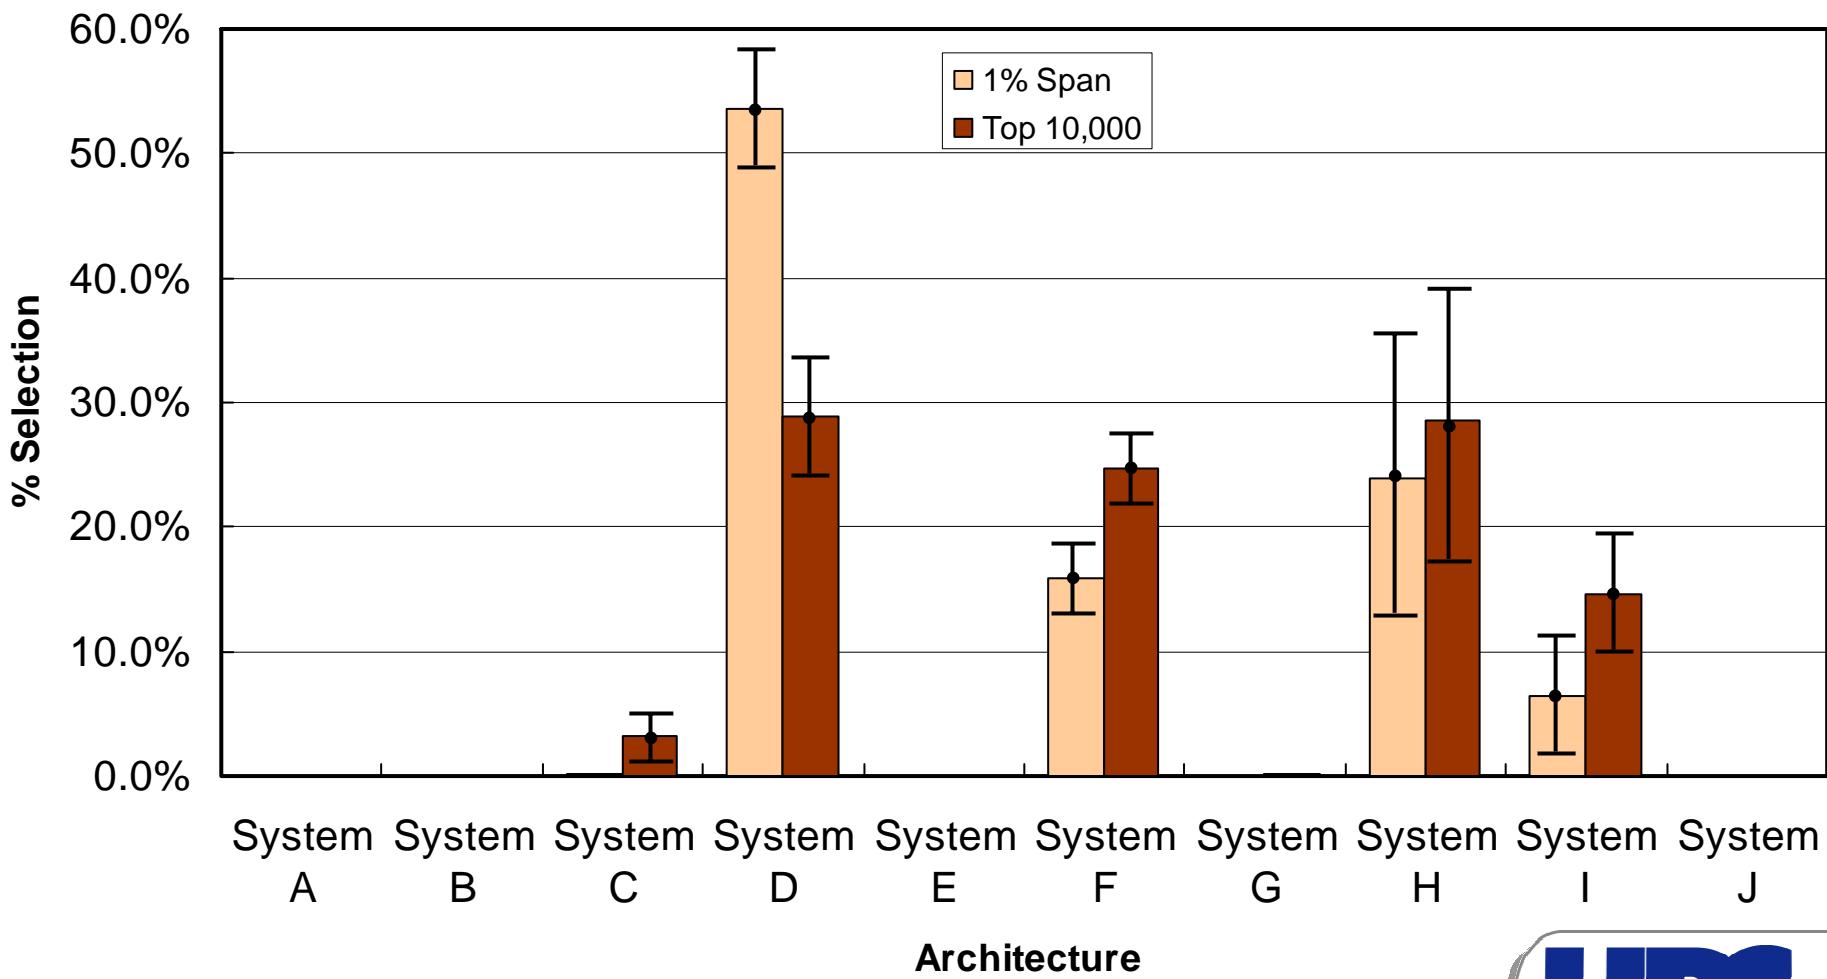


Propagation of Uncertainties in Benchmarking and Performance Modeling





U (EXIST+LC) Architecture % Selection by Processor Quantity for Varying Spans (TI-04)





Performance Measurement – Closing Thoughts

- Clearly identify your goals
 - Maximize the amount of work given fixed \$ and time.
 - Alternative goals: power consumption, weight, volume
- Define Work Flow
 - Production (run) time
 - Alternative goals: development time, problem set-up time, result analysis time
- Validate Measures
 - Understand the error bounds
- Don't rely on “Marketing” specifications!



2004 HPEC Conference

